

UNITED STATES PATENT APPLICATION
FOR

**METHOD OF DELIVERY, TARGETING, AND
MEASURING ADVERTISING OVER NETWORKS**

INVENTORS:
Dwight A. Merriman
Kevin J. O'Connor

PREPARED BY:

KENYON & KENYON
1500 K Street, N.W., Suite 700
Washington, DC 20005
(202) 220-4200

METHOD OF DELIVERY, TARGETING, AND MEASURING ADVERTISING OVER NETWORKS

Cross-Reference To Related Applications

[0001] This is a divisional of U.S. Patent Application No. 10/254,923, filed on September 26, 2002, pending, which is a continuation of U.S. Patent Application No. 09/293,463, filed on April 15, 1999, abandoned, which is a continuation of U.S. Patent Application No. 08/738,634, filed on October 29, 1996, issued as U.S. Patent No. 5,948,061, currently in reissue under Reissue Application No. 09/577,798, filed May 24, 2000.

Background Of The Invention

Area of the Art

[0002] This invention relates to methods of delivery of advertisements and measuring responses to those delivered advertisements and in particular relates to the targeting of advertisements delivered over networks such as the Internet.

Description of the Prior Art

[0003] In advertising, it is considered highly desirable to target advertisements to the appropriate potential customer base, rather than to broadcast advertisements in general. It has long been known that, for example, advertisements for computers should generally not appear in magazines on gardening and, conversely, advertisements for gardening tools should not appear in magazines on computers. Similarly, advertisers have generally targeted their advertisements on television to programs appropriate for the desired customer base.

[0004] It has also long been known that an advertisement that is repeated too often will eventually become ignored by consumers. Therefore, an advertiser typically wishes to eliminate duplication and reach as many individuals in the advertiser's target group as possible.

[0005] The recent development of on-line networks, such as America On-Line, Compuserve, and the Internet, has led to "on-line" advertising. For example, on the Internet, often such on-line advertisements will appear on a web page, such as a banner on the top or the bottom of the page. When the user views a web page using a browser such as Internet Explorer 3 or Netscape 3, the banner appears at the appropriate location and the user may then try to find out more information

regarding the advertisement by selecting the advertisement (clicking through on that banner) through the use of the mouse or other pointing device. This will cause a HTTP message to be generated by the browser using the information encapsulated in association with the banner to send a request for an object with a given URL address to a different appropriate web site to access, for example, the advertiser's home page.

[0006] Nonetheless, such advertising has had, so far, a poor rate of response because it is untargeted advertising. Thus, someone who is totally uninterested in computers other than they happen to be on the Internet, may continually see advertisements for computers. On the other hand, someone who is interested in computers may continually see advertisements for gardening tools when browsing through a particular web site. Thus it would be highly desirable to have a method of targeting the advertising to the appropriate user.

[0007] In addition, if a user of such computer networks is continuously exposed to the same advertisement, the response rate to the advertisement will generally decline. Therefore, it is highly desirable to have a system that controls the frequency of exposure of advertisements to particular users. In addition, it is also important for the advertisers to track response to the advertisements and to acquire as much information about those people responding to the advertisements for targeting those same people at later dates.

[0008] Therefore, it is a first object of this invention to provide targeting of advertising over networks such as the Internet. It is a second object of this invention to provide control over frequency of exposure to users for advertisements appearing on web pages over time. It is a third object of the invention to provide the capability to gather information about recipients of the advertisement.

Summary Of The Invention

[0009] These and other objects of the invention are achieved by the disclosed system and methods. Information about networks and subnetworks is routinely collected. In addition, information about individual users is also gathered when users select (click on) different advertisements. Also, data is tracked on how often a given advertisement has been displayed,

how often a given user has seen a given advertisement, and other information regarding the user and the frequency of the display of the advertisement

[0010] To effect such a capability, an advertising server process is provided as a node on the network. The various advertisements are stored on the network of the server and preferably on the server. When, for example, a user using a web browser accesses a web page that is affiliated with the advertising server process, the affiliated page's encoding includes an embedded reference to an object provided by the advertising server process. That causes the user's browser to contact the advertising server process to provide the advertising image or information that will appear on the accessed web page as displayed by the user's browser. Using the address information and/or other information passed by the browser for the user, including the page being accessed by the user, the advertising server process determines an appropriate advertisement to select for the particular user. In addition, the advertising server process will use information such as the number of times the user has seen various advertisements, how often the advertisement has been seen by any user and the start and stop date for the various advertisements to select which advertisement to transmit to the user's web page for display.

[0011] If the user decides to respond to the advertisement selected by the web server by clicking on the advertisement, the advertising server process logs that fact and to have more information about the given user. A derive profile process is used for compiling information on TCP/IP networks for use by the advertising server process. By compiling the information on networks and user selections, the advertising server process is able to compile information that can be used for targeting advertising.

Description of the Figures

[0012] FIGURE 1 is diagram for explaining a first embodiment of the invention.

[0013] FIGURE 2 is a diagram explaining the processes performed in the preferred embodiments.

[0014] FIGURES 3A, 3B and 3C are diagrams showing the basic structure of some of the databases kept by the advertising server.

Detailed Description of the Invention

[0015] The basic architecture of the network 10 comprises at least one affiliate web site 12, an advertisement (ad) server web site 19 and one or more individual advertiser's web sites 18.

Affiliates are one or more entities that generally for a fee contract with the entity providing the advertisement server permit third party advertisements to be displayed on their web sites. When a user using a browser accesses or "visits" a web site of an affiliate, an advertisement provided by the advertisement server 19 will be superimposed on the display of the affiliate's web page displayed by the user's browser. Examples of appropriate affiliates include locator services, service providers, and entities that have popular web sites such as museums, movie studios, etc.

[0016] The basic operation of the system is as follows in the preferred embodiment. When a user browsing on the Internet accesses an affiliate's web site 12, the user's browser generates an HTTP message 20 to get the information for the desired web page. The affiliate's web site in response to the message 20 transmits one or more messages back 22 containing the information to be displayed by the user's browser. In addition, an advertising server process 19 will provide additional information comprising one or more objects such as banner advertisements to be displayed with the information provided from the affiliate web site. Normally, the computers supporting the browser, the affiliate web site and the advertising server process will be at entirely different nodes on the Internet. Upon clicking through or otherwise selecting the advertisement object, which may be an image such as an advertisement banner, an icon, or a video or an audio clip, the browser ends up being connected to the advertiser's server or web site 18 for that advertisement object.

[0017] In FIGURE 1, a user operates a web browser, such as Netscape or Microsoft Internet Explorer, on a computer or PDA or other Internet capable device 16 to generate through the hypertext transfer protocol (HTTP) 14 a request 20 to any one of preferably a plurality of affiliate web sites 12. The affiliate web site sends one or more messages back 22 using the same protocol. Those messages 22 preferably contain all of the information available at the particular web site 12 for the requested page to be displayed by the user's browser 16 except for one or more advertising objects such as banner advertisements. These objects preferably do not reside on the affiliate's web server. Instead, the affiliate's web server sends back a link including an IP

address for a node running an advertiser server process 19 as well as information about the page on which the advertisement will be displayed. The link by way of example may be a hypertext markup language (HTML) tag, referring to, for example, an inline image such as a banner. The user's browser 16 then transmits a message 23 using the received IP address to access such an object indicated by the HTML tag from the advertisement server 19. Included in each message 23 typically to the advertising server 19 are: (i) the user's IP address, (ii) a cookie if the browser 16 is cookie enabled and stores cookie information, (iii) a substring key indicating the page in which the advertisement to be provided from the server is to be embedded, and (iv) MIME header information indicating the browser type and version, the operating system of the computer on which the browser is operating and the proxy server type. Upon receiving the request in the message 23, the advertising server process 19 determines which advertisement or other object to provide to user's browser and transmits the messages 24 containing the object such as a banner advertisement to the user's browser 16 using the HTTP protocol. Preferably contained within the HTTP message is a unique identifier for the advertiser's web page appropriate for the advertisement. That advertisement object is then displayed on the image created by the web user's browser as a composite of the received affiliate's web page plus the object transmitted back by the advertising web server.

[0018] As part of the "click through" process, when the user clicks on the banner or other advertising object displayed by the user's browser 16, the user's browser again transmits a message to the ad server. The ad server notes the address of the computer of the browser (or any other identifier such as a cookie or a digital signature) that generated the message 23 and transmits back the URL of the advertiser's web page so that the user's web browser 16 generates a message 26 to contact the advertiser's web site. 18. The ad server process 19 also notes that a "click through" for an advertisement has occurred and updates the various databases in the manner described below. In the above scenario for the click through process, the ad server process must remember which advertisement was sent to the user's browser in order to know where to redirect the user's browser.

[0019] While in the above embodiments, the user is a computer on an IP network using a browser, the affiliate web sites are web pages of affiliates located somewhere on the Internet and

the ad server is a particular node on the Internet, other setups are also possible. The affiliates may be ISP's or may be actual dedicated web servers and the users may be an entire network instead of an individual browsing on a single computer with a browser.

[0020] FIG. 2 shows the ad server architecture. The ad server, which may comprise one or more servers uses a database 54 that will be described below and performs reporting processes 59, management processes 58, derivation of profile processes 52 and advertisement processes 19. The derive profile process 52 is how the advertisement server gathers information about individual users or TCP/IP networks for individual users. Advertisements, which may be advertisement banners are stored within the ad server process 19 as part of the advertising server process 19 and are periodically updated and refreshed. The advertisement server process 19 is used for responding to requests from advertisements provided by the user's as described above. The management process 58 is used for updating the various advertisements and overall control of the advertising server process 19 and also permits the advertisers to interface with the database to obtain up to the date reports on the placement of the advertisements. The report process 59 is used for generating online reports about the success rate of the advertisement and statistics on the users that are viewing and clicking through on various advertisements and also updating the counters in the database that store how often an advertisement has been displayed.

[0021] The basic database structure is shown in FIGURE 3. For each user identified by the system as shown in FIGURE 3A, a user identification, IP address, domain type, time zone, location of the user, standard industrial code for the user's network, the particular advertisements seen and the number of times each such advertisement is seen, the advertisements that were selected or "clicked on" and the pages on which the various users' advertisements were seen is collected. By using the information such as which advertisements a user has expressed interest in and which pages the user was viewing when the user clicked through along with other information, targeted Internet advertising is possible.

[0022] FIGURE 3B shows a database structure for the advertisements. Included in the information for an advertisement are an identifier for the advertisement, the start date for the advertisement being carried, the last date the advertisement is to be carried, the total number of people who have viewed the advertisements, the target or the minimum number of times the

advertisement is to be viewed, acceptable viewers by their web site's SICS (Standard Industry Codes), clicked through and pages that the advertisements are seen on.

[0023] Also, as noted above, to enhance the process of identifying information about various users, information on domains or networks is also tracked. This information includes the domain name and type (e.g., military, government, commerce, foreign countries), the IP address, the standard industrial code, the time zone and the address as shown in FIGURE 3C. This information can be used in developing information about users. Through the use of these three databases and other appropriate databases that may be kept, the various processes of the advertising server process 19 are performed.

Advertising Server Processes 19

[0024] In response to an incoming message to the advertising server process 19, the advertising server processes 19 first attempts to identify the user in FIGURE 2. This can be accomplished by at least one of two means. First, any incoming request for an image or a multimedia object is examined to determine the IP address of the requesting browser. The advertising server then notes whether a cookie was received in the MIME heading of the request. From these two pieces of information, a user identification is determined. If a cookie was detected, then the cookie contains the user's identification number that can be accessed in the database. If the user's browser is cookie enabled but no cookie is detected, then the request is from a new user so a user identification must be assigned to the user and that user's new identification number will be transmitted back to the user's browser along with a write cookie instruction that causes the browser to write a cookie containing that unique identification number on the user's local drive for future accesses to affiliate sites. The cookie instruction is transmitted back with the advertisement messages 24. Alternatively, instead of using cookies, digital signatures or certificates or log ins uniquely identifying the user accessing the affiliate page may be used.

[0025] If the user's browser does not support cookies, the advertising server process looks up the user's IP address in a table that stores identification numbers that correspond to IP addresses for the users who have previously contacted the advertising server process. If the IP address is not found in the table, then the user is a new user and is assigned an unused identification number. Also, for each new user, that user's identification number will be marked for further processing

under the derive profile process 52 described below. Also, if the domain for the new user has not previously been processed in the domain profile process, it may not be possible to target the advertisement for the new user and rather the new user should be shown a generic advertisement. Also, for new users, a promotional advertisement may be shown to get the new user to provide information about him or herself and his or her employer.

[0026] If the user is an existing user, the ad server 19 obtains from a database all of the information known about the user including the user's geographic location, the domain type (commercial educational, governmental, the Internet service provided), the organization type where the user works (for example a SIC code), the company size, the number of employees in that company, the particular types of advertisements that the user has clicked on by SIC or other appropriate coding and the number of times that the user has been exposed to each advertisement currently in the system as described in FIGURE 3A. Also, the relative time of day for the user is calculated based upon either the user's country code or the user's IP access provider or the location of their domain.

[0027] Each advertisement along with a table of the targeting profile criteria for the advertisement and other data regarding the advertisement currently available is stored in a database such as shown in FIGURE 3B. The actual advertising object, which may be a banner image in a GIF or JPEG file format, an icon for an audio or video clip or some other object is kept as part of the advertising server process. This information may include targeted consumers by SIC, country, organization type and type of advertisements previously selected by the user. For any of the advertisements currently provided for which there is a match (for example there may be several advertisements targeted at people interested in computers), the next step is to determine which of the advertisements for which there is a match should be selected.

[0028] The preferred embodiment also includes determining which advertising object should be selected if two or more advertising object have criteria matching the user. selecting from the matched advertisements by determining how often the particular user has been exposed to the given advertisement. For each user, data is kept about the number of times that a user has seen a given advertisement based upon the user ID. For each advertisement where the user matches the criteria, if the number of times that the user has seen the advertisement is less than a

predetermined threshold, the advertisement is retained as one of the possible matches. If the number of times that the user has seen a given advertisement exceeds the threshold, the advertisement is discarded as a possible candidate for transmitting back to the user.

[0029] From the advertisements for which the user matches the criteria and for which the user has seen the advertisement less than the threshold, any of number of algorithms may be used to select the particular advertisement to transmit back to the user. In an embodiment, the formula used is based upon a satisfaction index (SI) according to the following formula:

$$SI = \frac{n}{N} * \frac{end - start}{now - start}$$

Where:

- n: the number of times the particular advertisement has been viewed by anyone
- N: the number of times the advertisement is to be see by anyone
- end-start: the total number of days that the advertisement is scheduled to run
- now-start: the number of days that the advertisement has run to date

[0030] The satisfaction index (SI) is 1.0 when the advertisement is on schedule. If the advertisement is running behind schedule (i.e., the current rate of display means that the advertisement will not be viewed the targeted number of times by the end date so that the SI is less than 1.0) while if the SI is greater than 1.0, the advertisement is running ahead of schedule. Therefore, the advertisement that is to be selected for viewing is the one that has a matched profile and has the smallest SI. Upon selecting the advertisement, the selected advertisement object is then transmitted to the user's browser over the Internet by transmitting the GIF stored in RAM in one or more messages.

[0031] After delivery of the advertisement to the user, additional processing needs to be done both for tracking the exposure of the advertisement and for having more information about the user. First, the system stores the fact that the advertisement was sent to the user by storing that information in the database based upon the user ID. This information is used for processing the user's response to the advertisement when the user "clicks" on the advertisement. Also, the

advertising processor updates the advertisements satisfaction index and logs the advertisement delivery for providing reports. Also, if the advertisement that has been sent has a limit on the number of times that the advertisement can be sent, the count of the number of times that the advertisement has been transmitted is incremented. This limit can be the overall limit for the number of times that the advertisement is viewed and/or the limit that any specific user can view the advertisement. If the limit on the number of times the advertisement can be sent is reached, then the particular advertisement is removed from the list of those that can be transmitted. In addition, even if an advertisement has not been transmitted recently, the SI is updated to reflect the passage of time.

[0032] If the user then decides that he wants to respond to the advertisement and then clicks on the advertisement, the advertisement server 19 receives an HTTP request 23 for a click through which it will recognize by receiving the same IP address (and optionally, the user ID specified by the cookie, digital signature or certificate or log in identification) in the request from the same affiliate's web page within a predetermined period. Since the advertisement server previously recorded which advertisement was sent to that IP address, it sends the redirect message causing the user's browser to receive the URL for the advertiser's web site based upon data stored in the server. In addition, the server logs that the advertisement was clicked through, which user selected the advertisement based upon the IP address, and the page on which the advertisement was seen based upon the click through. This information can then be logged in later reporting processes 59.

[0033] As an alternative to using a satisfaction index, the selection of the particular advertisement object to be sent back to the viewer's browser by the advertising server process 19 can be made by different criteria. For example, models can be created to maximize the likelihood of a click through. Such models for maximizing the likelihood of a click through may be based upon empirically gathered information about either types of advertisements or individual advertisements actually on the system correlated to the compiled information on a user or the user's network and the specific page through which the user is accessing the advertising server process 19.

Derivative Profile Process 52

[0034] In addition, as part of the profiling process for gathering information about users and IP addresses needs to be done. First, appropriate "networks" need to be identified based upon the TCP/IP network address of the user. Periodically, the derive profiles process searches the user database for the address of additional users that have addresses for unknown networks, indicating that they are new users. Then the server does a reverse domain name search to locate the domain name of the user. The server also queries Internet Whols databases to find information about the network, which will often include the domain name.

[0035] Alternatively, a reverse form of look up can be used independent of people accessing the network. When a domain is discovered, the server will check common DNS names for the name, such as those starting with "www" and "ftp". These resolve to IP addresses in most cases. From the IP addresses, the network number can be extracted and if the network does not yet have a domain name associated with the network number or address, the new domain associated with the network. A reverse domain name look up (A Whois lookup) will then usually provide the name, address and phone number of the organization, thereby providing the geographic location and the time zone of the network. Once the domain name is acquired, the server will determine whether the domain is an educational, military, governmental network and for non-U.S. based networks what country the network is located in through the extension.

[0036] However, it should be noted that some networks cover broad geographic areas and that further checking needs to be done to ensure that the data may be inaccurate. This can be done by performing a trace route operation to trace the network topology back to the network being examined. If the last two or three nodes on the trace show as being from a given region, the chances are high that the specific user on a network is located in that geographic area. Thus, reverse traces can be used to confirm location with a database for geographic locations. By checking telephone directories and other sources, the specific location for such users may be found by looking up the locations where the owner of the domain name has facilities and using the one closest to the outcome of the trace back step. If the information for a given user is different than the location of the domain based upon the trace back, it is this address and location information that is then used.

[0037] Further for each domain that is found, a determination must be made as to whether the domain is an Internet Service Provider (ISP). A list of domains that are ISP's is compiled manually. For any user having a domain name of an ISP, the user is presumed to be an individual and so only information relating to that individual are used for the profiling process.

Reporting Process

[0038] The reporting process uses standard database techniques for generating any desired reports. For example, reports showing the number of viewings and click throughs of various advertisements may be generated along with reports showing the click through rates relative to the matching criteria. Also, reports on the SI can be readily generated.

[0039] In addition to using the domain profiling process and logging advertisement clicked on by individual users and the pages being accessed by users of affiliate web sites, in addition surveying may be used. Further, for those survey respondents who are responding for non-ISP users discussed above, the results of the survey can be attributed to others on the same network. This provides further information for targeting the display of the advertisements.

[0040] Of course, those of skill in the field will understand that the disclosed techniques need not just be confined to Internet advertising but will work on other communication networks and private on line services such as CompuServe and America On Line. In addition, while specific user and domain profile information and matching criteria are discussed, it will be obvious to those of ordinary skill in the field that the specific type of user and domain name information profiled and used as matching criteria may include other or different criteria. In addition, while the specific examples are for IP networks, the same concepts can be applied to virtual LAN's such that a VLAN is the equivalent of a domain or on IPX based protocols or other network protocols.

[0041] Still further, rather than using cookies or IP addresses or other network addresses to identify the users, digital signatures or certificates may be used. In addition, a log in technique may be used whereby when users first access an affiliated web site, the user is assigned through a log in procedure a unique identifier and a log in name. When the user subsequently accesses either the same or a different affiliated web site, the user logs in using the unique identifier,

thereby ensuring that users can be identified notwithstanding proxy servers or different users using the same computer.

[0042] Still further, those of ordinary skill in the field will also understand that while the advertising server, the affiliate web site and the advertiser's web site are described as being in different geographic locations, that is not required. Still further, while the advertising server process, reporting process, derive profile process and management process are described as being implemented on one computer platform performing all of the above described functions, it is readily understood by those of skill that any or all of these functions may be implemented on one or more different computers and further that these processes may be performed at different nodes on the network.